

AMENDMENTS TO THE SPECIFICATION:

Please amend the title to read as follows:

METHOD, SYSTEM AND COMPUTER PROGRAM PRODUCT FOR MANAGING
DATA IN A MIRRORED CACHE USING AN ACCESS BALANCING TECHNIQUE.

Please amend the paragraph beginning at page 1, line 2 as follows:

Sub
A1 This application is a continuation in part of U.S. Patent application No. 09/676,686 filed
on September 29, 2000 (~~pending~~), now U.S. Patent No. 6,591,335.

Please amend the paragraph beginning at page 1, line 10 as follows:

A2 Host processor systems may store and retrieve data using a storage device containing a plurality of host interface units, disk drives, and disk interface units. Such storage devices are provided, for example, by EMC Corporation of Hopkinton, Hopkinton Mass. and disclosed in U.S. Patent No. 5,206,939 to Yanai et al., U.S. Patent No. 5,778,394 to Galtzur et al., U.S. Patent No. 5,845,147 to Vishlitzky Vishlitzky et al., and U.S. Patent No. 5,857,208 to Ofek. The host systems access the storage device through a plurality of channels provided therewith. Host systems provide data and access control information via the channels of the storage device and the storage device provides data to the host systems also through the channels. The host systems do not address the disk drives of the storage device directly, but rather, access what appears to the host systems as a plurality of logical disk units. The logical disk units may or may not correspond to the actual disk drives.

Please amend the paragraph beginning at page 2, line 19 as follows:

A3
U.S. Patent No. 5,437,022, U.S. Patent No. 5,640,530, and U.S. Patent No. 5,771,367, all to Beardsley et al, disclose a system having two, somewhat – independent, “clusters” that handle data storage. The clusters are disclosed as being designed to store the same data. Each of the clusters includes its disks own cache and non-volatile storage area. The cache from one of the clusters is backed up to the non-volatile data storage area of the other cluster and vice versa. In the event of a cache failure, the data stored in the corresponding non-volatile storage area (from the other cluster) is destaged to the appropriate disk. However, this system requires, in effect, a duplicate backup memory for each of the caches and also provides that whenever data is written to one of the caches, the same data needs to be written to the corresponding non-volatile storage in the other cluster. In addition, since each cluster includes a cache and a non-volatile storage, ~~thus~~ having two redundant clusters requires four memories (one cache for each of the clusters and one non-volatile storage for each of the clusters).

Please amend the paragraph beginning at page 3, line 14 as follows:

A4
In accordance with a first aspect of the invention, is a method of managing data in a cache is described. A first cache memory is provided that contains data. A second cache memory is provided that contains data, wherein at least some of the data in the first cache memory is the same as at least some of the data in the second cache memory. In response to a request for data that is stored in both the first cache memory and the second cache memory, one of the cache memories is chosen to use to obtain the data according to an access balancing technique.

Please amend the paragraph beginning at page 4, line 1 as follows:

A5 In accordance with another aspect of the invention, is a computer program product for managing data in a cache is described. Machine executable code is included for: providing a first cache memory that contains data; providing a second cache memory that contains data, wherein at least some of the data in the first cache memory is the same as at least some of the data in the second cache memory; and, in response to a request for data that is stored in both the first cache memory and the second cache memory, choosing one of the cache memories to use to obtain the data according to an access balancing technique.

Please amend the paragraph beginning at page 4, line 8 as follows:

A6 In accordance with yet another aspect of the invention, is a system for managing data in a cache is described. A first cache memory includes data. A second cache memory includes data, wherein at least some of the data included in the first cache memory is the same as at least some of the data of the second cache memory. Cache selection hardware is included for selecting, in response to a request for data that is stored in both the first cache memory and the second cache memory, which one of the first and second cache memories to use to obtain the data in accordance with an access balancing technique.

Please amend the paragraph beginning at page 4, line 17 as follows:

A7 Fig. 1A shows a pair of cache memories where each cache memory is coupled to a pair of buses in an embodiment of the system described herein.

Please amend the paragraph beginning at page 4, line 19 as follows:

A8

Fig. 1B shows a pair of cache memories coupled to a single bus in ~~an other~~ another embodiment of the system described herein.

Please amend the paragraph beginning at page 5, line 10 as follows:

A9

Fig. 6 is a flow chart illustrating steps performed in connection with a ~~host accessing data~~ in read operation executed by the host where data is being read from one or both of the cache memories.

Please amend the paragraphs beginning at page 6, line 1 as follows:

A10

Fig. 9 is a partial flow chart illustrating a portion of the steps performed in connection with a host accessing data in the cache memories including obtaining control data based on statistical analysis.

Fig. 10 is a partial flow chart illustrating a portion of the steps performed in connection with a host accessing data in the cache memories including obtaining control data based on a round robin scheme.

Fig. 11 is a partial flow chart illustrating a portion of the steps performed in connection with a host accessing data in the cache memories including reading data from another cache.

Fig. 12 is a partial flow chart illustrating a portion of the steps performed in connection with a host accessing data in the cache memories including read data based on statistical analysis.

A₁₀ Fig. 13 is a partial flow chart illustrating a portion of the steps performed in connection with a host accessing data in the cache memories including reading data from cache based on a round robin scheme.

Fig. 14 illustrates specialized hardware for providing the functionality illustrated in connection with Fig.'s Figs. 9-13.

Please amend the paragraph beginning on page 6, line 14 as follows:

A₁₁ Referring to Fig. 1A, a schematic diagram 20 shows a first cache memory 22, and a second cache memory 24 each coupled to a first bus 26 and a second bus 28. The cache memories 22, 24 and the buses 26, 28 may be part of a larger system, such as a data storage device provided by EMC Corporation of Hopkinton, Mass. Data may be written to and read from the memories 22, 24 via the busses 26, 28. The first memory 22 may be coupled to the first bus 26 via a first controller 32 and may be coupled to the second bus 28 via a second controller 34. Similarly, the second memory 24 may be coupled to the first bus 26 via a third controller 36 and may be coupled to the second bus via a fourth controller 38. The busses 26, 28 may be deemed "odd" and "even" for reference purposes. Similarly, the memories 22, 24 may be deemed "top" and "bottom".

Please amend the paragraph beginning on page 7, line 17 as follows:

A₁₂
Referring to Fig. 1B, a schematic diagram 30 shows an alternative embodiment where the first cache memory 22 and the second cache memory 24 are both coupled to a single bus 26'. In the embodiment of Fig. 1B, the bus 26' may be coupled to all of the host interface units and all of the disk controllers of the corresponding storage device. The system described herein may be configured with either the embodiment of Fig. 1A, the embodiment of Fig. 1B, or other configurations of one or more buses coupled to the cache memories 22, 24.

Please amend the paragraph beginning on page 8, line 1 as follows:

A₁₃
Referring to Fig. 2, a schematic diagram 40 illustrates a storage system 41 and the flow of data between the cache memories 22, 24, a disk storage area 42, and a host system 44. Data flows between the first cache memory 22 and the disk storage area 42 and flows between the first cache memory 22 and the host system 44. Similarly, data flows between the second cache memory 24 and the disk storage area 42 and between the second cache memory 24 and the host system 44. Specific control of the data between the ~~hosts~~ host system 44, the cache memories 22, 24, and the disk storage area is described elsewhere herein.

Please amend the paragraph beginning on page 8, line 8 as follows:

A₁₄
Referring to Fig. 3, a table 52, which is part of the data that is used to control operation of the storage ~~device~~ system 41, indicates portions T1, T2... TN of the cache memories 22, 24 that are to be designated as primary storage areas. In one embodiment, the cache memories 22, 24 are mapped alternatively so that, for example, a first set of portions may be designated as primary for the cache memory 22 while a second set of portions may be designated as primary for the cache memory 24, where the first and second sets are interposed. In some embodiments,

A14 the portions are 1/4 Gigabyte in size, although it will be apparent to one of ordinary skill in the art that the invention may be practiced using other sizes. The purpose of the mapping is discussed in more detail elsewhere herein.

Please amend the paragraph beginning on page 8, line 17 as follows:

A15 Referring to Fig. 4, a schematic diagram illustrates the cache memories 22, 24 in more detail. In some embodiments, each of the cache memories 22, 24 is implemented using separate hardware. Each of the cache memories 22, 24 is shown as containing a plurality of slots S1, S2, . . . SZ which, for embodiments discussed herein, provide storage for a sector of the disk storage area 42. For the embodiments illustrated herein, one sector equals eight blocks and one block equals five hundred and twelve bytes. However, it will be apparent to one of ordinary skill in the art that other sizes may be used without departing from the spirit and scope of the system described herein.

Please amend the paragraph beginning on page 9, line 10 as follows:

A16 Each of the slots represents data that is read from the disk storage area 42 and stored in one or both of the cache memories 22, 24. The control data for each of the slots indicates the state of the data in the slot. Thus, for example, the control data element for a slot can indicate that the data has been read from the disk storage area 42 but not written to by the host 44 (i.e., not modified by the host 44). Alternatively, the control data element for a slot could indicate that the data in the slot has been written to by the host 44 since being read from the disk storage area 42 (i.e., write pending). Note that, generally, data that is read from the disk storage area 42 but not subsequently modified may be eliminated from the cache without any ultimate loss of data

since the data in the cache memories 22, 24 is the same as the data in the disk storage area 42.

On the other hand, data that is write pending (i.e., modified while in the cache memories 22, 24

after being read from the disk storage area 42) ^{must be} is written back to the disk storage area 42 for

proper data synchronization. Note also that the control data could indicate that the associated

slot contains data that is the same in both of the ^{cache} memories 22, 24, which could occur, for

example, either when the data is write pending or immediately after data that is write pending is

written to the disk.

Please amend the paragraphs beginning at page 10, line 3 as follows:

In one embodiment, data that is read from the disk storage area 42 is written to one or the other of the cache memories 22, 24. The shading of the slots in the memories 22, 24 in Fig. 4

indicates that a slot has been designated as a secondary slot. Thus, for example, the slots S1, S2 .

. . SN of the cache memory 22 are designated as secondary slots while the slots SO, SP, . . . SQ

of the cache memory 24 are designated as secondary slots. Conversely, the slots SO, SP, . . . SQ

of the cache memory 22 are designated as primary slots while the slots S1, S2, . . . SN of the

cache memory 24 are designated as primary slots.

Please amend the paragraphs beginning at page 10, line 10 as follows:

In one embodiment, data that is read from the disk storage area 42 is written only to the corresponding primary slot and, at least initially, is not written to the secondary slot. Thus, for

example, if a sector of data is to be provided in slot S1, the data is read from the disk and is

initially written only to the cache memory 24. Similarly, data from the disk designated for slot

SP is initially written only to the cache memory 22. The hardware may be used, in a

conventional manner, to control writing to one of the cache memories 22, 24 or writing to both of the cache memories 22, 24 simultaneously (and/or with a single command). Similarly, the hardware may control which of the cache memories 22, 24 is read.

Please amend the paragraphs beginning at page 10, line 18 as follows:

If an event occurs causing data in the cache memories 22, 24 to change (such as a write from the host 44), then the modified data is written to both the primary memory and to the secondary memory. For example, data that is designated for slot S1 is initially written from the disk storage area 42 only to the cache memory 24. However, if a subsequent operation occurs that causes the data in slot S1 to change (i.e., a write by the host 44 to the ~~portion of the disk storage area 42 corresponding to slot S1~~ portion of the cache memory 24 corresponding to the disk storage area 42), then the data in slot S1 is modified according to the write operation which writes data to both of the memories 22, 24. Thus, data that is write pending exists in both of the cache memories 22, 24. Note that, in some instances, unmodified but related data in a slot may be copied from one of the memories 22, 24 to the other one of the memories 22, 24.

Please amend the paragraph beginning on page 11, line 16 as follows:

Note that any data that is write pending in the cache is provided in both of the cache memories 22, 24. On the other hand, data that does not need to be written back to the disks (i.e., data that has not been modified by the host 44) is stored in only one of the cache memories 22, 24. Storing the data in only one of the cache memories 22, 24 is an optimization that can increase performance by requiring only one write to one of the cache memories 22, 24 in certain instances, while also providing a mechanism where write pending cache data is written to both of

the cache memories 22, 24. In addition, note that, as discussed above, identical data may be stored in corresponding slots in both of the memories 22, 24 even though the data is not write pending. This may occur, for example, immediately after write pending data is copied to the disk.

Please amend the paragraph beginning on page 12, line 3 as follows:

Referring to Fig. 5, a flow chart 60 illustrates steps performed in the event that the hardware associated with one of the cache memories 22, 24 fails. Implementing each of the cache memories 22, 24 with separate hardware increases the likelihood that failure of the hardware for one of the cache memories 22, 24 will not occur at the same time as failure of the hardware for ~~an other~~ another one of the cache memories 22, 24. Detection of the failure of one of the cache memories 22, 24 is provided in a straightforward manner, such as described in U.S. Patent No. ~~5,742,501~~ 5,724,501 to Dewey et al., which is incorporated by reference herein. Note that detection of a failure may occur during an initial self test.

Please amend the paragraph beginning at page 12, lines 11 as follows:

Processing begins at a first step 62 where a pointer is set to point to the first slot of the good cache memory (i.e., the one of the cache memories 22, 24 that has not failed). Following the step 62 is a test step 64 where it is determined if the data stored in the slot that is pointed to is duplicated in the cache memories (i.e., is the same for both of the cache memories 22, 24). As discussed above, this is indicated by the corresponding control data for the slot. Note that this information is available irrespective of whether the slot of the non-failing one of the cache

A22
memories 22, 24 is a primary or a secondary storage area, since all of the control data is duplicated between the cache memories 22, 24, as discussed elsewhere herein.

Please amend the paragraph beginning at page 13, line 13 as follows:

A23
Note that the step 70 is also reached from the step 64 if it is determined that the data is the same in both of the cache memories 22, 24, and that the step 70 is also reached from the test step 66 if it is determined that the data, although not the same in both of the cache memories 22, 24, is stored in the non-failing one of the cache memories 22, 24. This is because, in either of these cases, it is not necessary to mark the control data for the slot being examined as indicating that the data is not in cache at the step 68.

Please amend the paragraph beginning at page 14, line 11 as follows:

A24
Following the step 82 is a step 84 where it is determined if the data is the same in both of the cache memories 22, 24. As discussed above, this information may be provided by the corresponding control data element. If it is determined at the test step 84 that the data is the same in both of the cache memories 22, 24, then the data may be read from either one of the cache memories 22, 24. Thus, if it is determined at the step 84 that the data is the same in both of the cache memories 22, 24, then control passes from the step 84 to a step 86, where the data is read from either of the cache memories 22, 24. In some embodiments, at the step 86 the data is read from the one of the cache memories 22, 24 that is used at the step 82 to obtain the control data. In other embodiments, at the step 86 the data is read from the one of the cache memories 22, 24 opposite to the one of the cache memories 22, 24 that is used at the step 82. Following the step 86, processing is complete.

Please amend the paragraph beginning at page 15, line 1 as follows:

A25
If it is determined at the test step 84 that the data ~~that~~ is not the same in both of the cache memories 22, 24, then control passes from the test step 84 to a test step 88 where the data is read from the primary cache for the data. The distinction between primary and secondary cache storage is discussed elsewhere herein. Following the step 88, processing is complete.

Please amend the paragraph beginning at page 15, line 5 as follows:

A26
Referring to Figure 7A, a flow chart 100 illustrates steps performed in connection with providing data from the disk storage area 42 to the cache memories 22, 24. At a first step 102, it is determined which of the cache memories 22, 24 is the primary storage area for the data by consulting the primary/secondary table. Following the step 102 is a step 104 where the data is copied from the disk storage area 42 to the one of the cache memories 22, 24 corresponding to the primary storage area. Following the step 104 is a step 106 where the corresponding control data element, for both of the cache memories 22, 24, is marked to indicate that the corresponding data is in cache, thus indicating that the data has been read ~~in to~~ into the cache. As discussed above, the control data for each of the slots of the cache memories 22, 24 is duplicated. Thus, the control data element for any slot in one of the cache memories 22, 24 is made to equal the control data for the slot in the other one of the cache memories 22, 24 by writing the control data to both of the memories 22, 24 at the step 106. Following the step 106, processing is complete.

Please amend the paragraph beginning at page 15, line 17 as follows:

A27
Referring to Fig. 7B, a flow chart 110 indicates steps performed in connection with the data in the cache that has been modified (e.g., by a write from the host 44). Note that the steps of

A27 the flow chart 110 may be executed some time after the data has been read from the disk storage area ~~41 in to~~ 42 into the cache or may never be executed at all for some of the cache data.

Please amend the paragraph beginning at page 16, line 5 as follows:

A28 Following the step 112 is a step 114 where the remainder of the sector that includes the modified block is copied from the primary cache to the secondary cache. As discussed above, the embodiments disclosed herein operate a sector at a time, although ~~is it~~ it would be apparent to one of ordinary skill in the art how to adapt the system to operate using different size data increments, such as a block. Thus, if the control data is provided on a per block basis, and if the cache holds and manipulates data in units of blocks, then it may be possible to forego the step 114. Note also that if the control data indicates that the data for the sector is the same in both of the cache memories 22, 24, then the step 114 may be omitted, since there would be no need to copy data that is already the same.

Please amend the paragraph beginning at page 16, line 15 as follows:

A29 Following the step 114 is a step 116 where the control data for the particular slot, in both of the cache memories 22, 24, is marked to indicate that the slot is write pending, indicating that the data has been modified while stored in the cache. As discussed above, the control data is written to both the primary and secondary storage areas. Following step 116, processing is complete. Note that when the write pending data is destaged, the control data may indicate that the data is no longer write pending although the control data may also indicate that the sector data in both of the cache memories 22, 24 is identical.

Please amend the paragraph beginning at page 17, line 12 as follows:

A30 Referring to Fig. 8, a flow chart 120 illustrates steps performed after the hardware for one of the cache memories has failed. Processing begins at a first step 122 which determines if the failed memory hardware has been replaced by a new memory board. The test step 122 represents waiting until new, operational, hardware for the failed memory board is installed. Thus, until the hardware for the failed memory is replaced, the step 122 loops back on itself. Stated differently, the remaining steps of the flowchart 120 are not performed unless and until the failed memory board is successfully replaced.

Please amend the paragraph beginning at page 17, line 19 as follows:

A31 Once the hardware for the failed memory has been replaced, control passes from the step 122 to a step 124 where the system is configured to write all data to both of the cache memories 22, 24. That is, every time data is read from the disk storage area 42 to ~~the~~ cache, or data that is in ~~the~~ cache is modified by the host 44, the data is written to both of the cache memories 22, 24.

Please amend the paragraph beginning at page 18, line 8 as follows:

A32 Following the step 126 is a test step 128 which determines if background copying is complete. If not, the step 128 loops back on itself to wait for completion. Otherwise, once background copying is complete, the cache memories 22, 24 are duplicates of each other and control passes from the step 128 to a step 130, where the system is reconfigured to operate in the usual manner as discussed above in connection with ~~the~~ Fig. 6, Fig. 7A, and Fig. 7B. Thus, when the hardware for one of the cache memories 22, 24 fails, the system operates with the single, non-failing cache memory. However, once the recovery process set forth in Fig. 8 is

A32 completed, then the system is reconfigured to have a primary and secondary cache and to operate in the usual manner, as discussed above.

Please amend the paragraph beginning at page 18, line 17 as follows:

A33 When the same data is stored ~~on~~ in both of the cache memories 22, 24, it is possible for the host system 44 to access the data from either of the cache memories 22, 24. Accordingly, in some instances, it may be possible to enhance performance by balancing access between the cache memories 22, 24. Depending upon the hardware configuration, it may be possible to access one of the cache memories 22, 24 while simultaneously accessing the other one of the cache memories 22, 24. Thus, balancing the accesses could enhance performance by increasing the number of simultaneous accesses and correspondingly decreasing the number of (inherently inefficient) serial accesses to the same one of the cache memories 22, 24.

Please amend the paragraph beginning at page 19, line 4 as follows:

A34 Referring to ~~Fig.'s~~ Figs. 9-13, a plurality of partial flow charts 80a - 80e illustrate modifications to the read operation illustrated in Fig. 6 and described above. The modifications facilitate balancing accesses of the data to increase throughput and enhance performance.

Please amend the paragraph beginning at page 19, line 7 as follows:

A35 The partial flow chart 80a of Fig. 9 illustrates an alternative step 82' that corresponds to the step 82 of Fig. 6 where the control data is accessed. As discussed above, for embodiments illustrated herein, the control data is the same for both of the cache memories 22, 24. The step

A35
82' represents choosing which of the cache memories 22, 24 to use for accessing the control data based on a statistical analysis. Although many different types of statistical analysis techniques could be used and would be apparent to one of ordinary skill in the art, for one embodiment of the invention, the statistical analysis of Fig. 82' simply tallies the number of accesses of each of the memories 22, 24 over a predetermined amount of time, such as one second. That is, if for the one second prior to the current access, the number of accesses of the cache memory 22 is N and the number of accesses to the cache memory 24 is M, then, if $N < M$, the cache memory 22 will be used, since the cache memory 22 is the one of the cache memories 22, 24 having the least number of previous accesses in the previous one second. Steps associate with the connector A' shown in Figure 9 are described in more detail elsewhere herein. 10

Please amend the paragraph beginning at page 20, line 6 as follows:

A36
Referring to Fig. 11, the partial flow chart 80c illustrates an alternative step 86' that corresponds to the step 86 of Fig. 6 where disk data that is the same in both of the cache memories 22, 24 is accessed. For the step 86', the disk data is read from the one of the cache memories 22, 24 that is opposite to the one of the cache memories 22, 24 used to access the control data. That is, irrespective of the technique for selecting which of the cache memories 22, 24 to use to access the control data, the step 86' represents reading the disk data from the other one of the cache memories 22, 24. Thus, for example, if the control data is read from the cache memory 22, then the step 86' represents reading the disk data from the cache memory 24. Similarly, if the control data is read from the cache memory 24, then the step ~~22~~ 86' represents reading the disk data from the cache memory 22.

Please amend the paragraph beginning at page 21, line 19 as follows:

A
37 Referring to Fig. 14, a diagram 150 illustrates specialized hardware 152 for providing the functionality, or at least a portion of the functionality, described above in connection with Fig.'s Figs. 9-13. The hardware 152 may be implemented using any one of a variety of technologies for designing customized hardware. The hardware 152 may be implemented using a single chip or a plurality of chips. The hardware 152 may receive access requests for the a cache via one or both of the buses 26, 28, or through (some means (not shown)). The hardware 152 may then process the requests in accordance with one or more of the techniques discussed above, and then read the data from one of the ~~eaches~~ cache memories 22, 24.

Please amend the paragraph beginning at page 22, line 5 as follows:

A
38 Note that using the hardware 152 may reduce the requirements of keeping additional statistics because the hardware may have direct access to the (a queue length) and other information used in connection with the techniques described herein.

Please amend the paragraph beginning at page 32, line 2 as follows:

A
39 Techniques A method, system, and computer program product are disclosed for managing data in a cache. A first cache memory is provided that ~~included~~ includes data. A second cache memory is provided that also includes data in which at least some of the data in the first cache memory is the same as at least some of the data in the second cache memory. In response to a request for data that is stored in both the first and second cache memories, one of the cache memories ~~are~~ is chosen in accordance with an access balancing technique. The access

A39 balancing technique may include ^{at least one of} selection using round robin, and selection based on statistical analysis such as access frequency of the first and second cache memories. 6
